

Management of OSI Layer-3 Data Network Entities

Field of the invention

[01] The invention relates to data network management, and in particular to methods of managing Open Systems Interconnection (OSI) Layer-3 data network entities.

Background of the invention

[02] Public data transport networks such as the Internet were initially designed to convey data encapsulated into what are known as Protocol Data Units (PDUs). Each PDU contains header information and a payload. The PDU represents a self contained entity holding routing information in the header to route the PDU in the data transport network towards a intended destination specified therein.

[03] The early development of the Internet was driven by a requirement of providing connectionless data transport. Connectionless data transport refers to the conveying of PDUs towards an intended destination data network node without specifying a route through the data transport network.

[04] The connectionless transport of data provides an inherent failure recovery. Each PDU is transported through the data transport network independently of any other PDU. The PDU is inspected at each data network node receiving thereof to determine a next hop to be used in forwarding the PDU towards the intended destination. The route taken by the PDU is sensitive to data link and/or data transport node failures. In performing routing decisions, the PDUs are routed around failures in the data transport network without necessitating redundant bandwidth and/or (physical) data links.

[05] Initially the Internet traffic primarily consisted of best effort data traffic by which it is meant that the data network nodes in the data transport network try their best to forward a PDU towards an intended destination without any guarantees of delivery. In cases in which a PDU is traversing a large number of data links, as well as in the case in which the PDU is awaiting processing at a congested data network node, it is possible for the PDU to be discarded or delayed. Data transport protocols, such as the Transmission Control Protocol (TCP), are used to provide a reliable data transport over unreliable data transport equipment. Although reliable, TCP data transport does not make Quality-of-Service (QoS) provisions: discarded or delayed PDUs are not acknowledged to have been received, after a delay unacknowledged PDUs are retransmitted.

[06] The paradigm of a data network path taken by PDUs as a “managed data network entity” does not make sense in this arrangement as PDUs are routed at every data network node and subsequent PDUs may not use the same data transport path in the data transport network by design.

[07] Recent developments in the data telecommunications field as well as recent market drivers have made increasing demands for data services requiring reliable data transport with a higher QoS. Equipment failures have been reduced by implementing redundancy. It was also realized that the data transport path taken by PDUs in providing data transport in support of a data service is very important.

[08] A large number of algorithms have been devised to use information about data network nodes and interconnecting data transport links to establish data transport paths having specific characteristics in supporting specific data services.

[09] Algorithms have been developed to find the shortest path between two data network nodes in the data transport network. By a shortest path one refers

to the shortest number of hops or physical data links to be traversed between source and destination data network nodes, or the shortest geographically traversed path between the data network nodes. These algorithms address the minimization of data processing delays and data transmission delays respectively. Other algorithms finding data transport paths attempt to optimize different aspects of data transport by assigning weights to data transport links as well as data transport nodes to optimize the chosen data transport path with respect to: data transport delay, data processing delays, jitter, bandwidth, redundant protection, etc.

[10] Examples of routing protocols implementing path discovery algorithms include, between others: Routing Information Protocol (RIP) distance based routing, Open the Shortest Path First (OSPF), Intermediate System to Intermediate System (ISIS), Border Gateway Protocol (BGP), etc.

[11] Connection oriented data transport technologies have also been devised to transport data in a fashion similar to circuit switching in which the establishment of a data transport path prior to data transmission involves at least some prior knowledge of the data transport network. Connection oriented data transport technologies include, but are not limited to the Multi-Protocol Label Switching (MPLS) protocol suite. MPLS embraces the "path as a data network entity" paradigm and particularly the "path as a manageable data network entity" paradigm.

[12] Routers represent another data network entity used in providing data transport services. Routers originate with repeaters. Repeaters were data network nodes used to extend the physical reach of data transport networks by blindly copying PDUs across the connected data network domains. Repeaters evolved into gateways as logic was added to only selectively copy the PDUs across if the intended destination data network nodes for the copied PDUs were not attached to the data network domain from which the PDUs were received.

Although gateways can be implemented as stand alone data network nodes, gateways are typically implemented integral to data switching nodes.

[13] With the advent of routing protocols implementing path discovery algorithms, as mentioned above, the gateways evolved into routers. Standalone
5 implementations of routers exist as well as virtual routers implemented in software running on data switching nodes.

[14] Provisioning data transport services ensuring Quality-of-Service, requires the establishment of data transport paths. Data transport paths are established between routers. Typical router management for standalone router
10 devices included the use of element management software. Element management techniques focuses on managing the router and the data transport paths associated with on a router by router basis. Performing router management using element management techniques suffers from a limited view of the data network resources. As opposed to data transport paths which
15 are governed by standardized data transport protocols, routers have a vendor specific implementation using vendor and typically router specific element management tools. Therefore network management solutions providing data transport services at QoS guarantees require the skills of an analyst having experience in configuring multi-vendor routing equipment.

[15] Therefore, methods are necessary for providing comprehensive management of data transport paths and routers in support of Quality-of-Service guarantees.

Summary of the invention

[16] In accordance with an aspect of the invention, a network management
25 system is provided. The network management system includes: network management logic; a human-machine interface; at least two view panes displayed on the human-machine interface. A first view pane displays

representations of OSI Layer 3 entities. A second view pane displays representations of OSI Layer 2 entities. The OSI Layer 3 entities are provisioned over the OSI Layer 2 entities.

[17] In accordance with another aspect of the invention, a network management software application is provided. The network management software application includes: network management logic; a human-machine interface for display of connectivity information, and at least two view panes displayable via the human-machine interface. A first view pane displays representations of OSI Layer 3 entities. A second view pane displays representations of OSI Layer 2 entities. The OSI Layer 3 entities are provisioned over the OSI Layer 2 entities.

[18] The combination of the displayed information in the at least two view panes provides for concise presentation of OSI Layer 2 entities corresponding to selected OSI Layer 3 entities. The advantages include an improved efficiency in the management, troubleshooting and provisioning of data transport paths.

Brief description of the drawings

[19] The features and advantages of the invention will become more apparent from the following detailed description of the preferred embodiments with reference to the attached diagrams wherein:

FIG. 1 is a schematic diagram showing OSI Layer-1 data transport elements as well as OSI Layer-3 data transport entities;

FIG. 2 is a schematic diagram showing a derivation hierarchy of managed entity objects modeling data transport network elements in accordance with an exemplary embodiment of the invention;

FIG. 3 is a schematic diagram showing a containment hierarchy of instantiated data transport managed entity objects in modeling the data

transport network elements in accordance with an exemplary embodiment of the invention;

FIG. 4 is a schematic diagram showing interconnected data transport network elements providing data transport for at least one data service provisioned via a data transport path;

FIG. 5 is a schematic diagram showing, in accordance with a preferred embodiment of the invention, associations between managed entity object instances used in provisioning data transport paths;

FIG. 6 is a schematic diagram showing, in accordance with an exemplary implementation of the invention, a graphical user interface enabling an analyst to interact with displayed graphical entities representative of instances of data transport managed entity objects;

FIG. 7 is a schematic diagram showing a progression in presenting connectivity information corresponding to a data transport path in accordance with another embodiment of the invention; and

FIG. 8 is a schematic diagram showing a progression of displayed information in troubleshooting a OSI Layer-3 data transport path to determine OSI Layer-2 and -1 failed equipment in accordance with yet another embodiment of the invention.

[20] It will be noted that in the attached diagrams like features bear similar labels.

Detailed description of the embodiments

[21] The Open Systems Interconnection (OSI) hierarchy specifies seven layers defining data transport functions performed by data transport protocols.

Isolating data transport functions in different layers minimizes the impact of technological change. Of relevant importance herein are the first three layers.

[22] OSI Layer-1 is also known as the Physical Layer concerning itself with physical data transport and in particular defines the characteristics of the hardware needed to carry out data transmissions. The hardware includes physical link cables, physical link optical fibers, free space physical links, etc. The physical link cables and/or physical link optical fibers connect to data ports. Each data port has physical transmitters and receivers including: signal drivers, laser modulators, microwave antennae, etc.

[23] OSI Layer-2 is also known as the Data Link Layer concerning itself with reliable data transport across corresponding Layer-1 physical links. Layer-2 data transport protocols include: Asynchronous Transmission Mode (ATM), Frame Relay (FR), Packet-over-SONET (Synchronous Optical NETWORK), Serial Line IP (SLIP), PPP, MAC, etc. Data network equipment includes nodes and interfaces having Layer-2 data transport protocol specific data ports. Data network ports may be configured to support more than one data transmission protocol either severally or in combination.

[24] OSI Layer-3 is also known as the Network Layer concerning itself with the management of data connectivity end-to-end. The connectionless Internet Protocol (IP) and the connection oriented Multi-Protocol Label Switching (MPLS) protocol are Layer-3 routing protocols. The MPLS protocol may make use of IP technologies. Data transport paths and routers are OSI Layer-3 data transport entities. MPLS enabled routers are known as Label Switching Routers (LSR).

[25] As mentioned above, for the original connectionless best effort data traffic, the idea of an ever changing data transport path as a manageable data network entity had little utility. Recently however, with the advent of connection oriented data transport technologies providing Quality-of-Service

guarantees make the “data transport path as a manageable data network entity” paradigm a requirement. With data transport paths spanned between routers in a data transport network, management of router entities is also a requirement.

5 [26] FIG. 1 is a schematic diagram showing data network elements implementing connected data transport networks.

[27] Data network nodes 102, 110, 120 are physically interconnected, to form data transport networks 100, at the Open Systems Interconnection (OSI) Layer-1 via physical links 108. Data transport networks 100 may be bridged via bridge
10 data network nodes 104 to enable data exchange therebetween. Connected data transport networks 100 can be grouped defining areas of focus and influence for the purposes of network management.

[28] Physical links 108 provide OSI Layer-1 connectivity between data network nodes 102/104/110/120 physically conveying data for OSI Layer-2
15 data links between nodes 102/110/120 end-to-end. A Layer-2 data link may be provisioned over at least one physical link 108 – a specific sequence of physical data links 108 making up an OSI Layer-3 path 128. Routers 130 are used to steer data traffic in the data transport network 100.

[29] For example, with respect to data switching nodes schematically shown
20 in Fig. 1, a vendor may choose an integral device 110 implementation having a data switching processor and a group of ports 112, while another vendor may chose a customizable data switching node 120 implementation including: a switching fabric, an equipment rack divided into shelves, each shelf 122 having slot connectors for connection to physical interface cards, each physical
25 interface card 124 having at least one port 112. Although, conceptually, the two the data switching nodes 110 and 120 provide the same data switching function, each implementation is adapted for a different environment: the former data switching node 110 is more adapted to enterprise solutions as a

private data network node, perhaps also adapted to enable access to public data services; while the latter data switching node 120 is better adapted for high data throughput in the core of public data transport networks. Typically the former (110) implements a small number of data transport protocols while for the latter (120), data transport protocols are implemented on interface cards 124 and/or ports 112 – providing for a flexible capabilities deployment thereof. Routers, such as schematically shown in FIG. 1 at 130 may represent integral data network devices. Virtual routers (130) are integrated into data network nodes 102/104/110/120.

[30] Network management and is typically performed with the aid of at least one Network Management System (NMS) 140 connected to at least one node 102 associated with a data transport network 100.

[31] FIG. 2 is a schematic diagram showing, in accordance with an exemplary embodiment of the invention, a derivation hierarchy of managed entity objects modeling data transport network entities.

[32] In support of management of the manageable data transport entities shown in Fig. 1, each manageable data transport entity is modeled in software via a corresponding managed entity objects using, for example, Object Oriented Programming techniques. Other modeling techniques may be used without departing from the spirit of the invention. The managed entity objects typically form a managed entity derivation hierarchy 200.

[33] FIG. 3 is a schematic diagram showing, in accordance with an exemplary embodiment of the invention, a containment hierarchy of instances of managed data transport entity objects in modeling the data transport elements used.

[34] In providing the network management function for data transport networks 100 in a realm of influence, data network management software applications are used to extract connectivity information from interconnected data network elements such as shown in Fig. 1. The network management

software application may further be used to specify connectivity information in setting-up data services for example. A data network state, including connectivity information, service provisioning information, operational state information, etc., is captured in instances of managed entity objects (200), the instances of managed entity objects forming a containment hierarchy 300.

[35] FIG. 4 is a schematic diagram showing connected data network elements implementing a data transport network 100 providing data transport for at least one data service 400 between a service provider node 412 and a service user node 422.

[36] The service provider node 412 and the service user node 422 are connected to the data transport network 100, and in particular to edge nodes 102, via uplinks 408. Data network nodes 102 within data network 100 are interconnected by a physical data links 108 to provide data transport therebetween.

[37] In order to provide Quality-of-Service in conveying data via the data transport network 100 between the service provider node 412 and the service user node 422, a data transport path 128 is provisioned between the edge nodes 102A and 102B. Data transport paths such as shown at 128 represent OSI Layer-3 entities and are supported by connection oriented data transport protocols.

[38] Each data transport path 128 traverses a group of nodes 102 in data transport network 100. PDUs sent via the data transport path 128, tunnel through the data transport network 100 along the established data transport path 128 between the end nodes 102 thereof. PDU processing efficiencies are derived from the tunneling thereof through the data transport network 100 reducing PDU processing overheads at data network nodes 102 along the data transport path 128. The routers 130 are used to steer data traffic associated with the data service 400 onto the data transport path 128.

[39] FIG. 5 is a schematic diagram showing, in accordance with a preferred embodiment of the invention, associations between managed entity object instances used in provisioning data transport paths.

[40] Irrespective of the implementation of a router 130 as a standalone device or integrated with a data transport node 102, each router 130 is associated with a data transport node 102. Routers 130 concern themselves with the processing of PDUs and in particular with the processing of routing information associated with the first three OSI layers stored in PDU headers.

[41] Management and provisioning data structures are defined at each OSI Layer. Such data structures include what are known as router interfaces. In order for a router 130 to steer data traffic, the router 130 uses router interfaces. At each layer connectivity is achieved via the router interfaces and the router interfaces form a hierarchy.

[42] A Layer-3 logical interface 152 is used for each data transport path 128 associated therewith. The Layer-3 logical interface 152 has an associated Layer-3 IP specification. The IP specification of the Layer-3 logical interface 152 is used by IP routing protocols, the MPLS protocol and static routing to specify connectivity between routers 130.

[43] Typically, each Layer-3 logical interface 152 may have one of two configurations, although the invention is not limited thereto. A first Layer-3 logical interface 152 configuration is said to be "numbered". A numbered Layer-3 logical interface 152, includes an IP ADDRESS (ADDR) in the IP specification thereof. The second Layer-3 logical interface 152 configuration is said to be "unnumbered". An unnumbered Layer-3 logical interface 152, does not include an IP ADDR in the IP specification thereof. The use of unnumbered configurations for Layer-3 logical interfaces 152 provides for the conservation of IP addresses which is a concern considering the increasing demand for data

transport services. A sub-layer interface is further used to specify Layer-2 and -1 connectivity.

[44] A Layer-2 logical interface 154, also referred to as a link layer-2 interface, is used as a sub-layer interface. Link layer interfaces 154 are tailored according to the data transport protocol used for each corresponding data link 162 used. Exemplary data transport protocols used include, but are not limited to: ATM, FR, etc. Corresponding layer-2 interface 154 specifications include Virtual Path Identifiers (VPI) / Virtual Circuit Identifiers (VCI) and Data Link Connection Identifier (DLCI) respectively. PPP session specification may also be used to configure the Layer-2 interface 154.

[45] A Layer-1 logical interface 156, also referred to as physical layer-2 interface, is also used as a sub-layer interface. Physical layer interfaces 156 are tailored according to the data transport used as well as the port 112 configuration used for each corresponding physical link 108 used. Exemplary data ports 112 used include, but are not limited to: ATM, FR, POS, Gigabit Ethernet, Ethernet (ports), etc.

[46] Data transport paths 128 include, but are not limited to: Layer-3 IP links provisioned using IP protocol, IP routing protocols and related technologies as well as Label Switched Paths (LSP) provisioned using the MPLS protocol. Ultimately the configuration of each IP link determines what types of PDUs can be conveyed over the IP link and therefore what data services can be provisioned. Exemplary IP link configurations provide support for IP forwarding, MPLS/IP forwarding, MPLS only, etc.

[47] The specification of an IP link is used to configure the above mentioned router interfaces and represents logical IP connectivity between routers 130 provisioning a data transport path 128. Only one IP link can be defined between a pair of router Layer-3 interfaces. One or more sub-layer interfaces (of the same type) may be associated with a router Layer-3 interface 152. More

than one IP link can be used between any two given routers 130 and each IP link may span more than one hop. An example of a multi-hop IP link may be provisioned over an ATM Virtual Non-Cross Connect layer-2 path. When using ATM as the Layer-2 data transport protocol, the IP link specification may

5 employ only one VPI/VCI pair. When using POS as the Layer-2 data transport protocol, the IP link may use multiple physical links in parallel with multiple (L1) ports being specified for each router interface (L3). An LSP may span over more than one IP link. LSPs may be limited to unidirectional directional data transport.

10 [48] Each IP link has two endpoints. For the purposes of network management, IP links fall into at least three categories: Point-to-Point, Point-to-IP, Point-to-Subnet, etc. In provisioning data transport paths 128, the first two IP links are used to establish point-to-point connectivity between routers 130, while the later is used provide support for broadcasting data

15 services. Each Point-to-Point IP link is used to configure router interfaces of both ends of the IP link associated with managed routers 130 in a realm of influence of the network management solution. If one of the routers 130 at an end of an IP link is not in the realm of management of the network management solution, then a Point-to-IP IP link is used only specifying the IP

20 ADDR of the unmanaged router 130.

[49] Further classification of IP links include: peer IP links used to establish connectivity between a router 130 within data transport network 100 to a bridge 104, a service provider (412) router or a service user (422) router; access

25 IP links used to establish connectivity between a service provider (412) router and a service user (422) router; and network IP links between managed routers 130.

[50] The Layer-3 IP link hides the layer-2 complexity as the Layer-3 logical connectivity is provisioned over a single or multi-hops as well as Layer-2 paths. The specification of an IP link does not configure IP routing protocols (when

enabled) but it does configure the MPLS protocol. Routing protocols used by routers 130 are configured separately and include but is not limited to the use of element management tools.

[51] The IP link/data transport path 128 establishment further includes the specification of QoS parameters such as, but not limited to: delay tolerance, jitter tolerance, PDU loss tolerance, forwarding priority, etc.

[52] FIG. 6 is a schematic diagram showing, in accordance with an exemplary embodiment of the invention, a graphical user interface enabling an analyst to interact with displayed graphical entities representative of instances of data transport managed entity objects.

[53] Preferably a display interface 600 is presented to the analyst. The display interface is divided in at least two view panes.

[54] A first view pane 610 shows at least a selected portion of the data transport network 100 in accordance to a Layer-3 representation thereof. Routers 130 steering data traffic onto the data transport path 128 are shown as well as the data path 128 itself. A pipe/tunnel spanned therebetween routers 130 represents the data transport path 128. Routers 130A and 130B may either be implemented as stand alone physical data transport equipment or implemented in software/firmware of a data transport node such as the virtual router 130B. In accordance to either implementation, the routers 130 are associated with data transport nodes 102 in the data transport network 100. Other nodes 102 bearing Layer-3 IP ADDRs are shown as part of the network management solution and further provide a visual context for the routers 130 and data transport path 128.

[55] The data transport network 100 itself is a Layer-3 entity to the extent that it has an IP network address ###.###.###.0 and may be displayed to provide a visual context.

[56] The data service 400 provisioned over the Layer-3 path 128 is not shown because it is a higher OSI layer entity. The service provider 412 and the service user 422 are not shown even though they are Layer-3 entities by bearing IP ADDRs because these nodes are not typically in the realm of influence of the network management solution.

[57] These Layer-3 representations besides showing the routers 130, the end nodes 102 and the data transport path 128 provide very little other information about. Although the information is limited, such representation is preferred eliminating unnecessarily complicating details associated with the lower OSI Layer-2 and -1 data transport equipment provisioning the data transport path 128.

[58] A second view pane 620 shows Layer-2 connectivity information associated with selected 604 Layer-3 entities in pane 610 – for example the data transport path 128. The selection is typically made via the use of pointing capabilities of the NMS 140 system.

[59] The Layer-2 connectivity information includes the link layer interface configuration specifications. Displayed in FIG. 6 is a multi-hop IP link providing data transport for data transport path 128. The multi-hop link specifies: the VPI/VCI for a hop 622 between ATM nodes 102A and 102D, the VPI/VCI for a hop 624 between ATM nodes 102D and 102C, and DLCI for a hop 626 between FR nodes 102C and 102B. It will be noted that node 102C is a gateway providing data transport protocol translation between ATM and FR. Optionally a data transport bandwidth allocated to the data transport path 128 may also be shown.

[60] FIG. 7 is a schematic diagram showing, in accordance with another exemplary embodiment of the invention, a progressive representation providing greater detail of connectivity information with respect to a provisioned path. For the purposes of this example, ATM data transport

technologies will be referred to in describing the relevant concepts of the invention.

[61] In managing data transport paths 128, an analyst (not shown) interacts with the NMS 140. In accordance with a high level Layer-3 representation of the connectivity information corresponding to a selected data transport path 128, the analyst is provided with a schematic representation thereof showing the data transport path 128 spanned between Routers 130A and 130B.

[62] In managing the data transport path 128, it is important to be able to inspect the underlying data network entities provisioning the data transport path 128. In accordance with the example shown, the data transport path 128 makes use of two data transport technologies and in particular ATM data transport 710 between end node 102A and data transport node 102C. Data network node 102C, a gateway, performs data transfer protocol conversion of the conveyed data to a FR data transport protocol for data transport 720 between node 102C and data network node 102B. This level of information includes OSI Layer-2 specific information relating to the use of the data transport protocols and protocol conversion at node 102C.

[63] In managing the selected data transport path 128, the analyst, by interacting with the NMS 140, may drill-down through the connectivity information associated therewith by interacting with the displayed iconical elements.

[64] For example, further drilling-down through the connectivity information associated with the data transport path 128, the analyst is able to ascertain that the end node 102A is an ATM aggregation node, that the ATM data transport 710 between nodes 102A and 102C has two hops via an intermediary ATM node 102D. Two data links 160 are shown: a first data link 712 between ATM nodes 102A and 102D, and a second data link 714 between ATM node 102D and the gateway node 102C. The Layer-2 representation includes

the Layer-2 configuration information including the VPI/VCI used over the data links 712 and 714.

[65] The data transport 720 between the gateway node 102C and the FR edge node 102B is shown to be provided via a FR data link spanned between the
5 ADM node 102C and the ATM edge node 102B. The Layer-2 representation includes the Layer-2 configuration information including the DLCI used over the data link 722.

[66] Further details of the physical Layer-1 devices used in provisioning the data transport path 128 may be presented including specifications of
10 shelves 122 used and slot specifications corresponding each physical L1 interface. It is important to note that the shelf and slot specifications may only be shown selectively as they are relevant only to specific data transport equipment.

[67] Further drilling-down through the connectivity information interface card 124 and physical port 112 specifications may be displayed. Again, the
15 interface card specification may only be shown selectively as it is relevant only to specific data transport equipment.

[68] FIG. 8 is a schematic diagram showing a progressive presentation of connectivity information with respect to data transport path 128 in
20 troubleshooting thereof to identify failed equipment. For the purposes of this example, SONET data transport technologies will be referred to in describing the relevant concepts of the invention.

[69] A Layer-3 data transport path 128 spanned between routers 130A and 130B is experiencing a failure. The NMS 140 may signify this fact to the
25 analyst via a visual queues such as special coloring, special purpose shading, special purpose animation, etc. In FIG. 8 the experienced failure is presented via heavy highlight. Upon selecting the data transport path 128, the Layer-2 connectivity information associated therewith is shown.

[70] By inspecting the presented Layer-2 connectivity information, the analyst determines that the failure is associated with SONET data transport between SONET node 102A and Add/Drop Multiplexer (ADM) node 102C over the POS data link 810. In drilling-down further through the Layer-2 connectivity information, the analyst is able to determine that the failure is associated with an OC-192 span 814 between ADM node 102D and the ADM node 102C.

[71] Drilling-down through Layer-1 connectivity information, a shelf 122C on the ADM node 102D and shelf 122D on the ADM node 102C are also shown to be affected by the failure, as well as corresponding slots. The information is presented for visual inspection by the analyst.

[72] Further drilling-down through Layer-1 connectivity information the physical interface cards associated with the physical link 814 namely 124C and 124D, as well as the data ports 112C/112D and corresponding channels.

[73] The troubleshooting session may include drilling-down through Layer-1 the connectivity information where a transceiver associated with an optical cable connected to the data port 112D is shown to experience a failure. The failure may be confirmed perhaps by failing a closed loop test performed on the transceiver.

[74] In accordance with the invention, the drill-down process zeroes-in more and more onto the failed equipment as well as eliminates from the view, properly functioning equipment even though associated with data transport path 128.

[75] The data transport path 128 specifications displayed, may include a data transport capacity and also provide further details with respect to a virtual tributary used over a particular SONET span.

[76] In accordance with the invention, the NMS 140 programmed with network management logic enables the analyst to navigate the containment hierarchy or instantiated managed entity objects modeling data transport equipment installed in the field. Selection routines programmed into the network management logic provide for the selective presentation of Layer-3 connectivity information on a human-machine interface 600 associated with the NMS 140. Further selection routines are invoked to progressively expose, and display lower layer connectivity information on demand. The display of connectivity information includes the display of Layer-3 connectivity information in the view pane 610 separate from the display of lower layer connectivity information in the view pane 620.

[77] The embodiments presented are exemplary only and persons skilled in the art would appreciate that variations to the above described embodiments may be made without departing from the spirit of the invention. The scope of the invention is solely defined by the appended claims.